Proceedings of the 2003 Winter Simulation Conference
S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, eds.

A MONTE CARLO SIMULATION APPROACH TO THE CAPACITATED MULTI-LOCATION TRANSSHIPMENT PROBLEM

Deniz Özdemir
Enver Yücesan

Technology Management Area
INSEAD
Boulevard de Constance
77305 Fontainebleau CEDEX, FRANCE

Yale T. Herer

Faculty of Industrial Engineering and Management
Technion
Israel Institute of Technology
Haifa 32000, ISRAEL

ABSTRACT

We consider a supply chain, which consists of N retailers and one supplier. The retailers may be coordinated through replenishment strategies and lateral transshipments, that is, movement of a product among the locations at the same echelon level. Transshipment quantities may be limited, however, due to the physical constraints of the transportation media or due to the reluctance of retailers to completely pool their stock with other retailers. We introduce a stochastic approximation algorithm to compute the order-up-to quantities using a sample-path-based optimization procedure. Given an order-up-to S policy, we determine an optimal transshipment policy, using an LP/Network flow framework. Such a numerical approach allows us to study systems with arbitrary complexity.

1 INTRODUCTION

Physical pooling of inventories has been widely used in practice to reduce cost and improve customer service. On the other hand, information pooling, which entails the sharing of inventory among stocking locations through lateral transshipments, has been less frequent. Transshipments, the monitored movement of material between locations at the same echelon, provide an effective mechanism for correcting discrepancies between the locations' observed demand and their available inventory. As a result, transshipments lead to cost reductions and improved service without increasing system-wide inventories.

Our research is motivated by observations from different industries. For example, transshipments are common in the management of spare parts. In manufacturing, factories turn to sister plants to quickly obtain a spare part before contacting the original supplier. Airlines have similar practices. Container shipping lines pool their containers through an exchange. Transshipments are increasingly common in apparel, fashion goods, and toys, particularly by those retailers with brick and click outlets.

All these transshipment practices, however, represent a *reactive* approach to unexpected stockouts. We believe that, if we take transshipment opportunities into account *proactively* during the planning phase, they can work as an effective mechanism for correcting demand-supply discrepancies, thereby reducing cost and improving service.

The literature on transshipments has generally addressed either problems with two retailers, e.g., Tagaras (1989), Tagaras and Cohen (1992), and Robinson (1990), or problems with multiple, identical retailers, e.g., Krishnan and Rao (1965) and Robinson (1990). In contrast, we consider multiple retailers, who may differ both in their cost structures and in their demand parameters. We further consider limits on transshipment quantities. Such capacity constraints may reflect the physical constraints of the transportation media or the reluctance of the retailers to completely share their stock with other retailers. Other recent work on transshipments includes Archibald et al. (1997), Tagaras (1999), Rudi et al. (2001), Herer and Rashit (1999), and Dong and Rudi (2000).

In this system, it is optimal for each retailer to follow an order-up-to S policy. The optimality of the order-up-to S policy takes into consideration the use of transshipments among retailers, to be performed once demand is observed. We show how the values of the order-up-to quantities can be calculated using a a stochastic approximation algorithm that is based on Infinitesimal Perturbation Analysis (IPA).

IPA has originally been introduced as a simulation-based optimization technique (Ho et al. 1979). With IPA, instead of using finite differences in a gradient search method, we use the mean value of the sample path derivative, which is obtained through a single simulation. In other words, we conduct a single simulation run, keeping track of the impact of a change in a system parameter value on performance. We then average these changes to estimate the gradient. The implicit assumption is that the average of these changes represents the change in expectations –hence, it yields an unbiased estimator. Glasserman (1991) established the general conditions for the unbiasedness of the IPA estimator. Applications of perturbation analysis have been reported in simulations of Markov chains (Glasserman 1992), inventory models (Fu 1994), manufacturing systems (Glasserman 1994), finance (Fu and Hu 1997), and control charts for statistical process control (Fu and Hu 1999). IPA-based methods have also been introduced to analyze supply chain problems (Glasserman and Tayur 1995).

Simulation-based derivative estimates help the search for an improved policy while allowing for complex features that are typically outside of the scope of analytical models. While the optimal order-up-to quantities have to be found once for the entire system, an optimal transshipment strategy has to be found on a period-by-period basis, given the period's demand realization. We show how these transshipment quantities can be found using an LP / Network flow framework.

The contribution of this paper is twofold. First is the development of an integrated IPA/LP algorithm for a system that allows capacitated transshipments. The system we consider is more general than previously studied systems with transshipments in that we consider multiple retailers, which differ both in their cost structure and in their demand parameters. Second is a methodological contribution, obtained by formulating and validating IPA derivative estimators for the transshipment problem. Formulating these estimators means introducing appropriate algorithms; validating them calls for showing that they converge to the correct values, where convergence is over the number of independent simulation replications used to estimate the derivative information.

2 PROBLEM DESCRIPTION

We consider a system with one supplier and N non-identical retailers, associated with N distinct stocking locations. The system inventory is reviewed periodically. The demand distribution of each retailer in a period is assumed to be known and stationary over time.

The first event in each period is the arrival of orders placed in the previous period. These orders are used to satisfy any outstanding backlog and to increase the inventory level. Next in the period is the occurrence of demand. Since the realization of demand represents the only uncertain event of the period, once it is observed all the decisions of the period, i.e., transshipment and replenishment quantities, are made. Lateral transshipments are then executed, and subsequently the demand is satisfied. Unsatisfied demand is backlogged. At this point, backlogs and inventories are observed, and penalty and holding costs, respectively, are incurred. The inventory is carried, as usual, to the next period.

Note that items in stock elsewhere in the system are supplied immediately through transshipments while backlogged items have to wait until the beginning of the next period. Thus, the advantage of using transshipments is in gaining a source of supply whose reaction time is shorter than that of the regular supply.

We now introduce the notation used. We will represent the vector of quantities described below, as well as the ones that we will introduce later in the paper, by dropping the subscripts, thus, $d = (d_1, ..., d_N)$.

- N = number of retailers;
- D_i = random variable associated with the periodic demand at retailer i with $E[D_i] < \infty$;
- $f(D)$ = joint probability density function for the demand vector D ;
- d_i = actual demand at retailer i in an arbitrary period;
- h_i = holding cost incurred at retailer i per unit held per period;
- p_i = penalty cost incurred at retailer i per unit backlogged per period;
- c_i = replenishment cost per unit at retailer i;
- $\hat{t}_{ij}$ = direct transshipment cost per unit transshipped from retailer i to retailer j;
- t_{ij} = effective transshipment cost, or simply the transshipment cost, per unit transshipped from retailer i to retailer j, $t_{ij} = \hat{t}_{ij} + c_i - c_j$.

3 OPTIMAL POLICIES

In each period, the replenishment and transshipment quantities must be determined. For the uncapacitated version of the problem, Herer et al. (1999) have already proven that, if transshipments are only made to satisfy the actual demand and not to build up inventory, there exists an optimal order-up-to $S = (S_1, S_2, ..., S_N)$ replenishment policy for all possible transshipment decisions. In the capacitated system, this property is preserved. We therefore adopt a base stock replenishment policy.

Given an order-up-to S policy for the replenishment quantities, the optimal transshipment quantities need to be determined each period between every two retailers. We define the decision variables in Table 1. In particular, let $F_{B_iM_j}$ represent the quantity transshipped from retailer i to retailer j. Recall that S_i is the order-up-to level at retailer i.

We also use the following auxiliary variable:

I_i = inventory level at retailer i immediately after transshipments and demand satisfaction

$$= S_i - \sum_{j=1}^{N} F_{B_iM_j} + \sum_{j=1}^{N} F_{B_jM_i} - d_i$$

Note that I_i may be either positive or negative, and we denote:

$$I_i^+ = \max\{I_i, 0\},\ I_i^- = \max\{-I_i, 0\}.$$

Now, the total cost of the system in a given period is given by:

$$TC = \sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} F_{B_i M_j} + \sum_{i=1}^{N} h_i I_i^+ + \sum_{i=1}^{N} p_i I_i^- \qquad (1)$$

In the above equation we have not fully accounted for the replenishment costs. Since we are using an "Order-up-to S_i" replenishment policy at each retailer, the total amount replenished system-wide will be exactly equal to the system-wide demand. Recall, however, that the replenishment cost differentials were included in the definition of t_{ij}. Thus, to fully account for the replenishment costs, we would need to add the term $\sum_{i=1}^{N} c_i d_i$ to Equation (1). However, since this term is independent of our decision variables, it is omitted.

Since the optimal policy is to order up to S_i units at retailer *i,* the point in time after an order arrives is a *regeneration point*. That is, the system returns to the same state (S_i units at each retailer *i* and no backorders) just after the start of each and every period. Thus, we can view the multi-period problem as a series of single-period problems.

Consider the movement of material in an arbitrary period. At the beginning of the period there are S_i units in stock at each retailer *i*. This stock can be used in one of three different ways: satisfy demand at retailer *i*, satisfy demand at retailer *j* (i.e., a transshipment from retailer *i* to *j*), and hold in inventory at retailer *i*. While it is true that it is physically possible to move stock from retailer *i* to another retailer, e.g., *j*, for storage, this is precluded.

At the end of the period replenishment arrives from the supplier. This material can be used in two different ways: to satisfy a backorder at a retailer or to buildup inventory at a retailer so that the period will end with the right amount of stock. The stock at the beginning of the period and the replenishment that arrives from the supplier are the only two sources of material during a period.

On the other hand, the demand at retailer *i*, d_i, can be satisfied in one of three different ways: from the initial inventory at retailer *i*, from the initial inventory at another retailer *j* (i.e., a transshipment from retailer *j* to retailer *i*), or from replenishment at the end of the period. Another sink for material is the requirement that each retailer *i* ends the period with S_i units in inventory. These units can come from one of two sources: the starting inventory at retailer *i* or replenishment at the end of the period. As discussed above, inventory from another retailer will not be used to buildup inventory levels at retailer *i*.

Using the observations above we model the movement of stock during a period as a network flow problem. In particular we have a source node, B_i, to represent the beginning inventory at retailer *i,* and a source node, R, to represent the replenishment that occurs at the end of the period. The sink node associated with the demand at retailer *i* will be denoted M_i and the sink node associated with the ending inventory at retailer *i* will be denoted E_i. The arcs in the network flow problem are exactly those activities described above and are summarized (with their associated cost per unit flow) in Table 1.

The complete network flow representation of the problem can be found in Figure 1 for three retailers. Note that the graph is bipartite, though our representation of the graph, which was chosen to show the connection to the underlying inventory problem, does not emphasize this point. The LP formulation associated with this network flow problem is as follows:

Problem (P)

$$Z(S,d) = \min \sum_{i=1}^{N} h_i F_{B_i E_i} + \sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} F_{B_i M_j} + \sum_{i=1}^{N} p_i F_{RM_i}$$

s.t.

$$S_i = F_{B_i M_i} + \sum_{\substack{j=1 \\ j \neq i}}^{N} F_{B_i M_j} + F_{B_i E_i} \quad i = 1,\ldots,N \qquad (2)$$

$$F_{B_i M_i} + \sum_{\substack{j=1 \\ j \neq i}}^{N} F_{B_j M_i} + F_{RM_i} = d_i \quad i = 1,\ldots,N \qquad (3)$$

$$\sum_{i=1}^{N} d_i = \sum_{i=1}^{N} F_{RM_i} + \sum_{i=1}^{N} F_{RE_i} \qquad (4)$$

$$F_{B_i E_i} + F_{RE_i} = S_i \quad i = 1,\ldots,N \qquad (5)$$

$$F_{B_i M_j} \le C_{ij}^{Tr}, \quad i, j = 1,2,\ldots,N, i \neq j \qquad (6)$$

$$F_{B_i E_i}, F_{B_i M_i}, F_{B_i M_j}, F_{RM_i}, F_{RE_i} \ge 0,\ i = 1,\ldots,N,\ j = 1,\ldots,N.$$

Equations (2), (3), (4) and (5), respectively, represent the inventory balance constraint at the B_i, M_i, R and E_i nodes. Equation (6) reflects a physical constraint, C_{ij}^{Tr}, on the quantity that can be transshipped from location *i* to location *j*. Alternatively, each location may wish to allocate only a portion, say β, of its on-hand inventory to trans-

Table 1: The Definition of the Arcs in the Network Flow Problem

Arc	**Variable**	**Cost per unit flow**	**Meaning**
(B_i, E_i)	$F_{B_iE_i}$	h_i	inventory is held at retailer *i*
(B_i, M_i)	$F_{B_iM_i}$	0	stock at retailer *i* is used to satisfy demand at retailer *i*
(B_i, M_j)	$F_{B_iM_j}$	t_{ij} $(t_{ii} = 0)$	stock at retailer *i* is used to satisfy demand at retailer *j*, i.e., transshipment from retailer *i* to retailer *j*
(R, M_i)	F_{RM_i}	p_i	shortage at retailer *i* is satisfied through replenishment
(R, E_i)	F_{RE_i}	0	inventory at retailer *i* is increased through replenishment

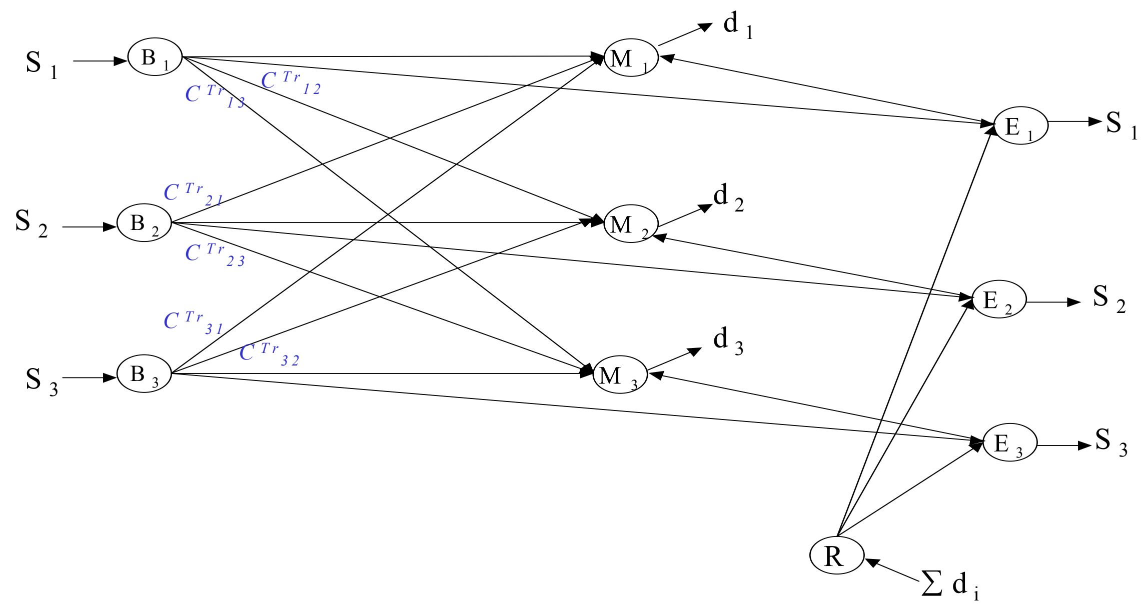

Figure 1: Network Flow Representation of a Single Period

shipments. This practice, typically referred to as *partial pooling*, can be represented through the following constraints:

$$\sum_{\substack{j=1 \\ i \neq j}}^{N} F_{B_iM_j} \leq \beta S_i, \quad i = 1,\ldots,N. \tag{7}$$

3.1 Finding the Optimal Order-up-To Values

In the most general setting, exact computation of optimal order-up-to levels by analytical methods is difficult. This is in fact the problem of optimizing an expected value function. Since the corresponding expectation function cannot be computed exactly, it is approximated through Monte Carlo sampling. Using the notation of Shapiro (2001), this represents a class of optimization problems of the form:

$$\min_{x \in X} \{g(x) := E[G(x, \omega)]\}, \tag{8}$$

where the expectation $g(x)$ is well defined for every $x \in X$. The function $G(x, \omega)$ is in itself an optimization problem. In our case, G is the optimal value of a network flow problem , where retailer demand is the random data of the problem.

We solve the optimization problem (8) by Monte Carlo simulation, that is, by generating an IID random sample and calculating the corresponding sample average:

$$\hat{g}_U(x) := U^{-1} \sum_{j=1}^{U} G(x, \omega^j).$$

The optimization problem (8) is then approximated by:

$$\min_{x \in X} \hat{g}_U(x). \tag{9}$$

We propose an IPA-based approach to solve (9). The idea is to use the expected value of the sample path derivative obtained via simulation instead of using the derivative of the expected cost in a gradient search method. In other words, the gradient of interest is dE[TC]/dS whereas our numerical procedure computes E[dTC/dS]. To validate this approach, that is, to justify the interchange of the derivative and the integral, we need to show that the objective function is jointly convex and "smooth" in the S_i variables (Glasserman and Tayur 1995). For the capacitated transshipment problem, this is done in Özdemir et al. (2003). We now turn to the description of the IPA algorithm.

3.2 Description of the IPA Algorithm

The algorithm starts with an arbitrary value for the order-up-to levels, S. An instance of the demand is generated at each retailer. Once the demand is observed, problem (P) is solved in a deterministic fashion to compute the minimum-cost solution. The gradient of the total cost (derivatives with respect to the order-up-to levels) is estimated and accumulated through independent replications; the average gradient value is then used to update the values of S. The procedure is summarized in a pseudo-code format, where K denotes the number of steps taken in a path search, U represents the number of independent replications, a_k represents the step size at iteration k, and S_i^k represents the order-up-to level for retailer i at the k^{th} iteration.

```
Initialize K
Initialize U
Set k ← 1
For each retailer, set initial order-up-to
levels, S_i^0
Repeat
Set dTC ← 0
Set u ← 0
Repeat
i. Generate an instance of the demand at each
   retailer, d, from f(D)
ii. Solve problem (P) to determine optimal
   transshipment quantities
iii. Accumulate the desired gradients (de-
   rivatives) of the total cost, dTC
iv. u ← u + 1
Until u = U
v. Calculate the desired gradient(s), dTC / U
vi.Update the order-up-to-levels, S_i:
    S_i^k ← S_i^(k-1) + a_k(dTC_i/U)
vii. k ← k + 1
Until k = K
```

In step (iii) of the algorithm, we use IPA to compute the gradient. To illustrate the sample-path derivative idea, suppose that we end a period with inventory at retailer i. In this case, raising S_i by 1 unit would result in increasing total cost by h_i. In the computer implementation, for each retailer i, we could partially code Step (iii) as:

$dTC_i = dTC_i + h_i$, if inventory at retailer i is positive, at the end of Step (ii).

Starting with dTC_i = 0 for all i at the beginning of the simulation and dividing dTC_i by U in Step (v) yield the derivative estimates.

Our network flow formulation greatly simplifies computations. Increasing S_i corresponds to increasing the supply at source node B_i and the demand at sink node E_i. From a network flow perspective, dTC/dS_i = h_i, if the arc (B_i, E_i) is basic or, equivalently, the flow $F_{B_iE_i}$ is positive. If the arc is non-basic, then since any basic solution corresponds to a spanning tree in the network, there exists a unique augmenting path from B_i to E_i whose total cost yields the gradient value. For example, the augmenting path may go from B_i to M_j to R to E_i, with an associated cost of $t_{ij} - p_j$. Such a path represents a transshipment from retailer i to retailer j (with a cost of t_{ij}), a reduction in backorders at retailer j (with a savings of p_j) and a purchase of another unit at retailer i (cost of zero).

Furthermore, our implementation of the derivative computation in Step (iii) is very efficient. Since the value of the gradient is equal to the total cost along the unique path from B_i to E_i for each retailer i, this quantity can be calculated directly as the difference between the holding cost at retailer i and the reduced cost of the arc (B_i,E_i), which is readily available from the linear programming solution in Step (ii).

In Step (vi) of the algorithm, one typically imposes conditions on the step size, a_k such that

$$\sum_{k=1}^{\infty} a_k = \infty \text{ and } \sum_{k=1}^{\infty} a_k^2 < \infty .$$

The first condition facilitates convergence by ensuring that the steps do not become too small too fast. However, if the algorithm is to converge, the step sizes must eventually become small, as ensured by the second condition. Note that when the gradient estimator is unbiased (as is the case here), step (vi) represents a Robbins-Monro algorithm (1951) for stochastic search.

4 COMPUTATIONAL STUDY

4.1 Experimental Design

An illustrative example of the system with four retailers is shown in Figure 2. Let us call retailer 0 the *central retailer* and all other N retailers the *remote retailers.* We begin by considering the case of identical retailers, the cost parameters are as follows: $h_i \equiv h$ = \$1 per unit, $p_i \equiv p$ = \$4 per unit, and the basic direct transshipment cost, c_t = \$0.5 per unit, when transshipments are allowed. Each retailer faces an independent demand stream distributed uniformly over (0, 200).

Note that t_{0i}, i=1,2,…,N, represents the transshipment cost from the central retailer to remote retailers, t_{i0}, i=1,2,…,N, represents the transshipment cost from the re-

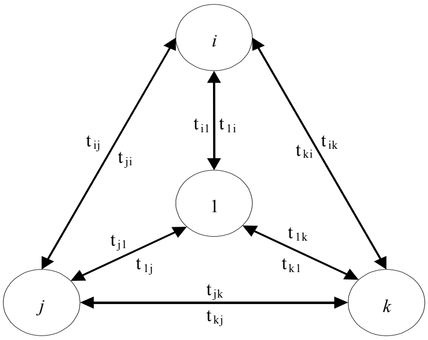

Figure 2: Configuration Used in Numerical Testing

mote retailers to the central retailer, and t_{ij}, i,j=1,2,…,N, denotes the transshipment cost from remote retailer i to remote retailer j. As summarized in Table 2, we consider five alternative system configurations. Note that $t_{ij} = \infty$ implies that transshipments are not allowed between retailers i and j.

Table 2: System Configurations

System	t_{0i}	t_{i0}	t_{ij}
1	∞	∞	∞
2	c_t	∞	∞
3	c_t	c_t	∞
4	c_t	c_t	$2\ c_t$
5	c_t	c_t	c_t

System 1, where no material movement is allowed among retailers, represents N+1 independent newsvendor problems. It thus serves as a benchmark. In system 2, transshipments are allowed only from the central retailer to the remote retailers. System 3 extends the scenario in system 2 by allowing transshipments from the remote retailers to the central retailer as well. In system 4, all material movement is possible. However, transshipments between any two remote retailers are twice as expensive as the transshipments from/to the central retailer. Finally, all transshipment costs are identical in system 5.

4.2 Results and Analysis

In all systems, we observe an increase in the total cost and in the total inventory levels when capacity considerations are incorporated. However, as the number of units allocated for transshipment increases, both the total cost and the total inventory levels decrease.

Figure 3 depicts the total cost as a function of transshipment capacity for the five systems with 10 retailers. In systems 2 and 3, where transshipment opportunities are restricted, decreasing capacity leads to increased costs. In systems 4 and 5, however, a tightening transshipment capacity is fully compensated for through the availability of transshipment opportunities between any pair of retailers.

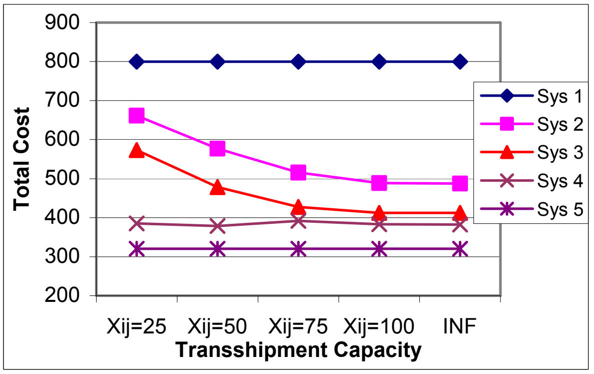

Figure 3: Total Cost in the Presence of Transport Constraints

Figure 4 illustrates an interesting redistribution of inventory in system 3 in the presence of transshipment capacities. Recall that, in system 3, stocking location 1 behaves as a clearinghouse for all other stocking locations. In other words, when there is no constraint on the quantity that can be transshipped, most of the material is stocked in location 1 and shipped to the other locations, as needed. In fact, up to 25% of the system-wide inventory is kept at location 1. However, when transshipment quantities are tightly constrained, location 1 finds itself unable to support any of the other locations through transshipments. It therefore does not have any reason to carry extra inventory. In fact, in a tightly capacitated environment, the other locations carry additional stock for location 1. Even for very small transshipment capacities, there is enough stock (9 times the transshipment capacity) carried by the other locations that can be sent to location 1 in case of a shortfall there.

5 SUMMARY

In this paper, we considered the multi-location dynamic transshipment problem, where transshipment quantities may be restricted. Our approach includes several innovations. First, an arbitrary number of non-identical retailers is considered with possibly dependent stochastic demand. Second, we model the dynamic behavior of the system in an arbitrary period as a network flow problem. Finally, we employ a simulation-based method using infinitesimal perturbation analysis for optimization. Our simulation-based optimization approach therefore provides a flexible platform to analyze transshipment problems of arbitrary complexity.

ACKNOWLEDGMENTS

This work has been partially supported by PriceWaterhouseCooper's grant to INSEAD's initiative on High-Performance Organizations.

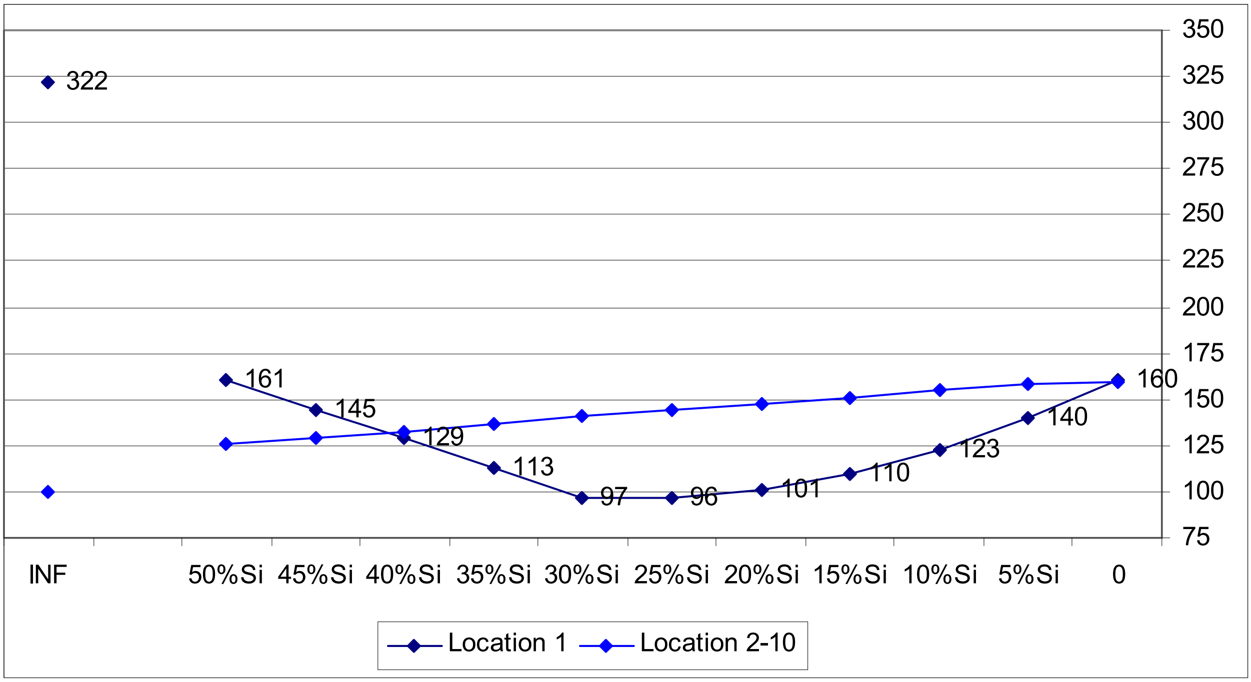

Figure 4: Optimal Base Stock Levels in the Presence of Transport Constraints

REFERENCES

Archibald, T.W., A.A.E. Sassen, and L.C. Thomas. 1997. An Optimal Policy for a Two-Depot Inventory problem with Stock Transfer. *Management Science* 43: 173-183.

Dong, L., and N. Rudi. 2000. Supply Chain Interaction Under Transshipments. Working paper. The Simon School, University of Rochester, Rochester, NY, U.S.A.

Fu, M.C. 1994. Sample Path Derivatives for (s,S) Inventory Systems. *Operations Research* 42: 351-364.

Fu, M.C., and Hu, J.Q. 1997. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Kluwer Academic Publishers.

Fu, M.C., and Hu, J.Q. 1999. Efficient Design and Sensitivity Analysis of Control Charts using Monte Carlo Simulation. *Management Science* 45: 395-413.

Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers.

Glasserman, P. 1992. Derivative Estimates from Simulation of Continuous-Time Markov Chains. *Operations Research* 40: 292-308.

Glasserman, P. 1994. Perturbation Analysis of Production Networks. In *Stochastic Modeling and Analysis of Manufacturing Systems* (Yao, Ed.). Springer-Verlag.

Glasserman, P. and S. Tayur. 1995. Sensitivity Analysis for Base Stock Levels in Multi-Echelon Production-Inventory Systems. *Management Science* 41: 263-281.

Herer, Y.T. and A. Rashit. 1999. Lateral Stock Transshipments in a Two-Location Inventory System with Fixed and Joint Replenishment Costs. *Naval Research Logistics* 46: 525-547.

Herer Y. T., M. Tzur, E. Yücesan. 2001. The Multi-Location Transshipment Problem. Working paper. Department of Industrial Engineering, Tel Aviv University.

Ho, Y.C., M.A. Eyler, and T.T. Chien. 1979. A Gradient Technique for General Buffer Storage Design in a Serial Production Line. *International Journal of Production Research* 17: 557-580.

Krishnan, K.S. and V. R. K. Rao. 1965. Inventory Control in *N* Warehouses. *Journal of Industrial Engineering* 16**:** 212-215.

Özdemir, D., E. Yücesan, and Y.T. Herer. 2003. Multi-Location Transshipment Problem with Capacitated Transportation. Working paper. Technology Management Area, INSEAD.

Robbins H. and S. Monro. 1951. A Stochastic Approximation Method. *Annals of Mathematical Statistics* 22: 400-407.

Robinson, L.W. 1990. Optimal and Approximate Policies in Multiperiod, Multilocation Inventory Models with Transshipments. *Operations Research* 38: 278-295.

Rudi, N., S. Kapur, and D. Pyke. 2001. A Two-location Inventory Model with Transshipment and Local Decision Making. *Management Science* 47: 1668-1680.

Shapiro, A. 2001. Monte Carlo Simulation Approach to Stochastic Programming. *Proceeedings of the 2001*

Winter Simulation Conference (Peters, Smith, Medeiros, and Rohrer, eds.) 428-431.

Tagaras, G. 1989. Effects of Pooling on the Optimization and Service Levels of Two-Location Inventory Systems. *IIE Transactions* 21: 250-257.

Tagaras, G. 1999. Pooling in Multi-Location Periodic Inventory Distribution Systems. *Omega* 39-59.

Tagaras, G. and M. Cohen. 1992. Pooling in Two-Location Inventory Systems with Non-Negligible Replenishment Lead Times. *Management Science* 38: 1067-1083.

AUTHOR BIOGRAPHIES

DENIZ ÖZDEMIR is a doctoral student in the Technology Management Area at the European Institute of Business Administration, INSEAD, in Fontainebleau, France. She can be contacted at <deniz.ozdemir@insead.edu>.

ENVER YÜCESAN is a Professor of Operations Management at the European Institute of Business Administration, INSEAD, in Fontainebleau, France. He has an undergraduate degree in Industrial Engineering from Purdue University. His MS and PhD, both in Operations Research, are from Cornell University. He served as the Proceedings Co-Editor for WSC'02. He can be contacted by e-mail at <enver.yucesan@insead.edu>.

YALE T. HERER joined the Faculty of IndustrialEngineering and Management at the Technion --- Israel Institute of Technology in 1990 immediately after the completion of his graduate studies at Cornell University and, except for a few year hiatus at the Department of Industrial Engineering at Tel-Aviv University, he has been there ever since. Presently Yale is serving as the Head of the Industrial Engineering Area. He has worked for several industrial concerns, both as a consultant and as an advisor to project groups. In 1996 Yale received the *IIE Transactions* best paper award. Yale is a member of the Institute for Operations Research and Management Sciences (INFORMS), Institute of Industrial Engineers (IIE), and the Operations Research Society of Israel (ORSIS). He serves on the Editorial Board of *IIE Transactions*. His research interests include supply chain management, especially when integrated with transshipments. He is also interested in production control and production system design. He can be contacted by e-mail at <yale@technion.ac.il>.